

subnetworks. Area 122 is used to define interconnections. Area 124 is used to provide details on objects and to allow a user to easily select objects for use.

Analysis

[88] Analysis includes monitoring and administration functions. Users can view results of node data-gathering which indicates the performance of system components, transfers, etc. Various administrative functions can be performed such as saving and modifying configurations, scheduling events,

[89] Four consoles, or basic types of interfaces, are used to help direct network optimization and manage the administration. The consoles are as follows:

1. Node Listing Console
2. Graphic View Console
3. Monitor Console
4. History Monitor Console

[90] Fig. 4A illustrates the Node Listing console.

[91] The Node Listing console provides a list of all the network nodes that are part of the current loaded network configuration, as well as the current status of those nodes. The console is also the location from which user access can be managed; different network configurations can be saved and loaded; backups can be initiated, and Wizards, or automated assistance, for redirectors and System Level Objects (SLOs) can be started.

[92] Fig. 4B illustrates the Graphic View console.

[93] The Graphic View console allows users to visually identify and manipulate the various nodes, pools and connections in a DASPO network in an easy-to-use graphical user interface.

[94] Fig. 4C illustrates the Monitor console. The Monitor console is a real-time tracking feature that measures the available processing capacity of selected nodes in DASPO network to help assess node performance. The node information is displayed in a simple graph or bar format, and the data can be tracked and saved for future reference.

[95] The Monitor console can provide several different graphs for visual presentation of information.

[96] Fig. 4D illustrates a series graph of the Monitor Console.

[97] In the series graph, selected SLO and TLO nodes appear with statistical values from 0 to 100 for each node at a given instant in time. The statistical value reflects the current

load capacity of the node. The higher the value, the more processing capability is available to be utilized. A lower value indicates an overworked node that has a low processing capacity.

[98] Host nodes that are selected to be monitored will appear in the Host graph. This graph performs identically to the Series graph.

[99] The Percentage graph measures the statistic values of SLO, TLO and Host nodes together on the same graph. This graph performs similarly to the Series and the Host graphs.

[100] Fig. 4E illustrates a balance graph of the Monitor Console.

[101] In the balance graph, statistical differences between the nodes is shown. Examples of types of differences that can be displayed include average, variance, maximum, minimum etc. These variances are shown visually on one or more bar graphs. A list of available balance variables can be selected and applied by a user. This graph appears beneath the Series and the Host graph in the iSystem Enterprise monitor. Note that the Balance graph does not appear when a Mixed Series is selected.

[102] Before node statistics or balance variables can be displayed in the Monitor graphs, the nodes to be monitored must first be selected. There are two selector fields at the bottom of the Monitor screen shown in Fig. 4E. The left-hand selector field 132 is used for adding SLO, TLO or Host nodes. The right-hand selector field 134 is used to add balance variables. (Note: the balance variable selector is not available when a Mixed Series is selected).

[103] Fig. 4F illustrates the History Monitor.

[104] When network nodes are tracked using the Monitor feature the captured data is stored, for future reference, in a log file. This log file can be accessed and displayed at any time using the History Monitor console. The History Monitor also provides a variety of features that allows saved data to be manipulated, displayed and compared in a variety of different ways. Note: In order to use the History Monitor feature, nodes must first be set up and tracked using the Monitor. For more information, see Monitor Console.

[105] The History Monitor provides several graphs similar to those described, above, for the Monitor Console.

[106] The History Monitor includes a series graph where monitored SLO and TLO nodes appear. This graph displays a statistical value (from 0 to 100) for each selected network node at a given instant in time. This statistical value reflects the load capacity of the node. The higher the value, the more processing capability is available to be utilized. A lower value indicates an overworked node that has a low processing capacity.

[107] Monitored Host nodes will appear in the Host graph of the History Monitor. This graph performs identically to the Series graph.

[108] The Percentage graph of the History Monitor displays the monitored statistic values of SLO, TLO and Host nodes together on the same graph. This graph performs identically to the Series and the Host graphs.

[109] The statistical differences between the nodes (i.e. average, variance, maximum, minimum etc.) can be measured in the balance graph of the History Monitor. A list of available balance variables can be selected and applied by a user. This graph appears beneath the Series and the Host graph in the iSystem Enterprise monitor. Note that the Balance graph does not appear when a Mixed Series is selected.

[110] Before the node statistics that have been captured in the monitor can be displayed in the History Monitor graphs, the nodes to be monitored must first be selected. There are two selector fields at the bottom of the History Monitor screen of Fig. 4F. The left-hand selector field 136 is used for adding SLO, TLO or Host nodes. The right-hand selector field 138 is used to add balance variables. (Note: the balance variable selector is not available when a Mixed Series is selected).

Optimization

[111] Part of the optimization process is accomplished by redirecting requests and connections within Functional Resource Pools. This is achieved using data generated by SLO-nodes, which compute their own statistics and broadcast the results through the pools.

[112] This way of implementing redirection is available to every application implemented in-house. However, there are many pre-packaged applications and objects commonly used, whose code cannot--and probably shouldn't--be altered. These types of applications include web servers and COM-objects. Due to the different nature of requests and connections that take place in a complex network system, specific objects must handle redirection inside each class of calls. A preferred embodiment of the present invention includes objects for redirecting HTTP-requests and OLE DB-connections. However, other embodiments can employ other objects in other environments and on other platforms such as HTTP in Java, DB in C++, etc., on Linux, Solaris, etc.

[113] An HTTP Redirector is a Windows-based application (HTTPRedir.EXE) capable of receiving HTTP-requests and redirecting them to a selected web server according to some predefined selection criteria. Starting from a list of web servers and a selection